

AD-A201 531

LARGE-SCALE, EXPLICIT WAVE SIMULATIONS ON THE CRAY-2

G.L. WOJCIK, D.K. VAUGHAN, M. BARENBERG and J. MOULD
Weidlinger Associates, 620 Hansen Way, Palo Alto, CA 94304, U.S.A.

M.B. HULIT

Air Force Geophysics Laboratory, Hanscom Air Force Base, MA 01731, U.S.A.

Most time-domain, wave modeling problems in geophysics are intractable by classical analysis methods, due principally to nonseparability and to a lesser extent material nonlinearity. Therefore discrete numerical solutions are often necessary for the simulation of realistic models. Applications in 2-D and 3-D geophysical modeling are the subject of this paper, particularly as solved on a CRAY-2 supercomputer. Implementation and performance differences between earlier CRAYs and the CRAY-2 are described, including the discrepancy between scalar fetch and vector processing speeds. Explicit finite element solvers are applied to applications involving 2-D simulation of a seismic refraction experiment across the state of Maine, 3-D simulation of near-source scattering experiments, and both linear and nonlinear axisymmetric source simulation. Results show that the CRAY-2 allows cost-effective 2-D simulations of truly large-scale models, but only begins to be effective in 3-D for models of interest in geophysics. The large memory (256 megawords) is adequate but a speed increase of at least an order of magnitude is necessary for cost-effective 3-D. True multiprocessor capability (rather than 'multi-computer') provides an alternative to raw speed but involves another set of difficulties. (F. F. F.). (FR)

1. Introduction

A large number of the time-domain, wave modeling problems in geophysics are intractable by classical analysis methods—by virtue of either nonseparability or nonlinearity. In fact, only a few practical problems are addressed by classical analysis, i.e., separable and linear, although this restricted class has received most of the theoretical attention. The broader class of nonseparable, nonlinear problems requires discrete numerical solutions. Some of these discrete time-domain wave propagation problems in geophysics are the subject of this paper, particularly as they are implemented and solved on the CRAY-2 supercomputer. The paper emphasizes various performance and modeling issues, with little attention given to analytical development.

By way of background recall that separability depends on the medium's degree of inhomogeneity, namely, whether it conforms to a separable coordinate system for the governing partial differential equations. Separability is not a good global assumption in general geophysical applications. In contrast, the typical assumption of linearity depends on the medium's constitutive model and, excluding the immediate source region, is often a good global assumption.

Separable problems include the common homogeneous or vertically stratified half-space and are well solved by classical methods, with weak inhomogeneities sometimes included via perturbation methods. For truly nonseparable problems, characterized by significant deviation from the layered half-space, no corresponding analytical methods are available. Either local

geometrical solutions or global numerical solutions are necessary, provided in practice by geometrical diffraction and ray theory or discrete numerical methods like finite difference and finite element wave solvers, and to a lesser extent by boundary integral methods.

Nonlinear continuum mechanics problems typically involve irreversible material behavior due to strains beyond the elastic limit (constitutive). Nonlinearities may also be due to large displacements and/or rotations (geometric). However, the contribution of geometric nonlinearities is generally secondary to constitutive effects in geophysical modeling. An effective means of including these constitutive nonlinearities is through a rigorous plasticity formulation.

The 'best' approach for practical global solutions of nonlinear, nonseparable wave propagation problems in geophysics is a discrete method. Reasons for this choice include ease of modeling, minimal need for geometric or material idealization, full wave representation, and the availability of efficient algorithms in conjunction with supercomputers. Certainly, discrete methods have their share of drawbacks as well, including the inability to generalize results of one calculation beyond its basic parameters, the difficulty of separating wave phenomena, e.g., body waves and surface waves, and errors associated with a finite model boundaries.

To address both the pros and cons of discrete numerical methods in geophysics, this paper describes some applications of explicit finite element solvers to large-scale wave modeling problems, principally on the CRAY-2 supercomputer. These include 2-D modeling of a refraction experiment across the state of Maine, 3-D modeling of some near-source scattering experiments, and both linear and nonlinear source modeling simulations.

2. Background

The discrete wave solvers applied here incorporate finite element reductions of the governing partial differential equations to ordinary differential equations in time. These ODEs are integrated forward in time using a modified leapfrog scheme (centered differences). The routines are included in a pre- and post-processing shell called FLEX, designed for efficient modeling, solution, and interpretation of large-scale propagation problems (mechanical or electromagnetic) [1]. The code was originally written to take advantage of the architecture and power of CRAY computers by developing fully vectorized kernel processing loops first, followed by an efficient code architecture to support them. Rather than use one general purpose processing routine, a group of specialized finite element and finite difference routines was developed for each class of problems, including 1-D, plane and axisymmetric 2-D, and full 3-D. FLEX was designed to minimize storage requirements for each problem by using an internal data management system that avoids the use of dimensioned arrays within the code. The system automatically and adaptively sets up internal data arrays as required by the problem at hand. Very large problems are efficiently stored in memory, thus avoiding I/O limitations of data transfer to disk.

2.1. Finite element solvers

The explicit finite element routines applied in this paper to linear and nonlinear continuum mechanics problems employ bilinear (2-D) or trilinear (3-D) shape or interpolation functions over each element, e.g. see [2]. Element geometry is either Cartesian (rectangles and bricks in 2-D and 3-D, respectively) or skewed. The term, explicit (in contrast to implicit), effectively means

or

☒

☐

☐

on

on/

Availability Codes

Dist	Avail and/or Special
A-1	20



that lumped masses are employed. This eliminates the need for assembly of a global stiffness matrix and a mass matrix inversion. Instead, nodal forces are accumulated element-by-element and the nodal (lumped mass) velocity increments are calculated from Newton's law in incremental form. The algorithm is stable provided the timestep is less than the fastest wave's transit time across the smallest element. This follows since all nodes are effectively decoupled by the fundamental hyperbolicity of the governing equations; in other words, the influence of a node during one timestep cannot affect nodes more than one element away. Numerical noise (roundoff) propagates one element per timestep.

The algorithm stores three types of quantities for each node—lumped mass $M(n)$ velocity vector $V(n)$, and force vector $F(n)$, where $n = 1, \dots, N$, with N the total number of nodes. The nodal velocity and force vectors' dimension, d , is equal to the problem dimension, so there are seven quantities per node in 3-D and five in 2-D (plane or axisymmetric). Velocities are defined at full timesteps while forces are defined at timestep midpoints. Nodal displacements ($U(n)$) defined at timestep midpoints are not calculated directly, but may be found by integrating velocities. Similarly, stresses are not explicitly calculated, but if required, e.g., in nonlinear calculations, they are defined at the element centroid at timestep midpoints and updated using incremental displacements.

The basic algorithm includes a vectorized force loop over rows of physically contiguous elements (row by row) and a velocity loop over all of the nodes. In the force loop, increments ($\Delta F = K\Delta U = KV\Delta t$) are calculated at the nodes of each element and added to the nodal force vector to obtain forces at the next half timestep ($F = F + \Delta F$). After all of the elements are processed a single nodal loop updates velocities. In this loop, velocity increments ($\Delta V = M^{-1}F\Delta t$, where M is diagonal) are calculated using the nodal forces at the previous half timestep, and added to the nodal velocity vector to obtain values at the next full timestep ($V = V + \Delta V$). These loops are repeated for the required number of timesteps.

Ninety percent of the floating-point operations in the algorithm occur in the element force loop. For example, in a 2-D plane, elastic model, this loop fetches nodal velocities and forces as well as material and shape information, performs 73 floating-point operations, and stores updated nodal forces for each element in the row. In general, data arrays for the nodal and elemental quantities are mapped into memory so that contiguous elements in a row have contiguous storage locations. The resulting data structures are such that a row of elements of any length can be integrated one timestep in a vectorized computation without the need of a gather operation to fetch data into a contiguous local array, and a scatter operation on the results.

The code lends itself to efficient vectorization in 'vanilla' FORTRAN (i.e., no assembly coding necessary), and approaches the peak performance levels expected for 'nonideal' problems, i.e., inhomogeneous, on pipelined supercomputers like the CRAY machines. This is not to say that a supercomputer is required for practical calculations: tens of thousands of elements are routinely executed on minicomputers, up to perhaps 100,000 elements, with reasonable execution times.

2.2. Nonlinear constitutive behavior

Nonlinear material behavior can be modeled by a variety of constitutive relations and numerical implementations. For soil-type and rock-type continua the so-called cap model has proven effective, particularly in the context of explicit finite element or finite difference wave

solvers. The cap model is essentially an algorithm implementing a rate-independent plasticity theory with an associated flow rule. Since the formulation is not as well known as the discrete wave propagation algorithms utilized in this paper, it is reviewed below.

The cap model, e.g., [3], is based on classical plasticity theory with incremental stress-strain relations in the form

$$\dot{\sigma} = C\dot{\epsilon}, \quad (2.1)$$

where σ is the stress tensor (tension positive), ϵ the strain tensor, and C the (tangent) modulus matrix, assumed positive semi-definite to insure uniqueness, stability, and continuity. The cap model is defined by a convex yield surface, $Y(\sigma)$, and a plastic strain rate vector, $\dot{\epsilon}^P$, which is normal to the yield surface in stress space so that

$$\dot{\epsilon}_{ij}^P = \lambda \frac{\partial Y}{\partial \sigma_{ij}}, \quad (2.2)$$

where λ is a non-negative scalar function.

Bulk modulus K and shear modulus G are used to represent simple linear elastic behavior inside the yield surface. The surface is defined by a fixed failure envelope and a hardening cap. The failure envelope is defined by

$$Y(\sigma) = \sqrt{J_2'} - F_F(J_1), \quad (2.3)$$

$$F_F(J_1) = A - C e^{BJ_1}, \quad (2.4)$$

where A , B , and C are material parameters, J_1 is the first invariant of the stress tensor, and J_2' the second variant of the deviatoric stress tensor. The cap is defined by

$$Y(\sigma) = \sqrt{J_2'} - F_c(J_1, \kappa) \quad \text{for } L(\kappa) \geq J_1 \geq X(\kappa), \quad (2.5)$$

where κ is an internal state variable that measures hardening as a functional of plastic volumetric strain history, and $L(\kappa)$, $X(\kappa)$ define the J_1 -range of the cap. Function $F_c(J_1, \kappa)$ is defined by

$$F_c(J_1, \kappa) = \frac{1}{R} \sqrt{[X(\kappa) - L(\kappa)]^2 - [J_1 - L(\kappa)]^2}, \quad (2.6)$$

in which

$$L(\kappa) = \begin{cases} \kappa & \text{if } \kappa < 0, \\ 0 & \text{if } \kappa > 0, \end{cases} \quad (2.7)$$

$$X(\kappa) = \kappa - RF_F(\kappa), \quad (2.8)$$

where R is a material parameter.

Hardening parameter κ is implicitly defined as a functional of the plastic volumetric strain, ϵ_v^P , by means of a relation between $X(k)$ and ϵ_v^P , which is then coupled with (2.8) to define k in terms of ϵ_v^P . For soils, the relationship is

$$\bar{\epsilon}_v^P = W[e^{D(X(\kappa) - X_0)} - 1], \quad (2.9)$$

in which W , D , and X_0 are material parameters and ϵ_v^p is a history-dependent functional of the plastic volumetric strain and is given by the differential functional relation

$$\dot{\epsilon}_v^p = \begin{cases} \dot{\epsilon}_v^p & \text{if } \dot{\epsilon}_v^p \leq 0 \text{ or } \kappa < 0, \\ 0 & \text{if } \dot{\epsilon}_v^p > 0 \text{ and } \kappa \geq 0. \end{cases} \quad (2.10)$$

If dilatancy occurs when $J_1 > 0$, (2.10) limits the shrinking of the cap to $\kappa = L(\kappa) = 0$. This insures that the cap remains finite and is assumed to apply in general for soils. Such a limitation is somewhat arbitrary in view of the lack of data with respect to material behavior after the occurrence of tension failure.

Tensile stresses are limited by the condition that the first invariant of stress

$$J_1 \leq T, \quad (2.11)$$

so that T (a material parameter) represents a tension cutoff parameter. In order to insure that no principal stress exceeds parameter T , the slope of the failure envelope (a measure of the friction angle) has been appropriately limited.

The cap model provides a theoretically sound idealization for the salient features of nonlinear soil behavior. The solution algorithm [3] is robust and quite efficient relative to many other nonlinear material models and has proven effective for a wide variety of applications.

3. CRAY-2 implementation

The CRAY-2 was the largest and fastest 'conventional' supercomputer available at the time the calculations described in this paper were performed (mid 1987). It uses a UNIX-based operating system, has four available central processing units, and addresses 256 million words of CMOS memory (typically). Other machines are available with theoretically faster architectures, e.g., the Connection Machine, but they also require significant modification of the discrete algorithms and are presently unproven for production problems.

The CRAY-2 is therefore the machine of choice for large-scale simulations, particularly since similar CRAY-1 type machines have been available for years and optimal programming techniques are fairly well understood. Note that the principal advantage of the CRAY-2 over earlier CRAYs (the X-MP for example) is the large memory available. The factor of two to three in speed is also significant but not the principal reason for its choice. The availability of four processors is not very useful in practice since chargeable time is accumulated on each one, with nontrivial overhead for multiprocessor usage. The only advantage is a factor of three or so decrease in wall clock time. This is useful for reducing run time, given the possibility of a system crash during extended runs. A better solution is to store memory images periodically for restarts during long runs (> five hours or so).

3.1. FLEX on the CRAY-2

Several differences were found between the performance of FLEX on the CRAY-2 and its performance on earlier CRAY machines. First, the original CRAY-1 FORTRAN coding produced erroneous results on the CRAY-2 caused by the use of CDIR\$ IVDEP system

commands that force vectorization of certain computational loops containing vector dependencies. Specifically, this occurs for the elemental loops computing internal force contributions to each node of an element. Since a node receives force contributions from two adjacent elements, it is clear that if the internal force contributions of both elements are being computed during the same pass through the loop, there is a potential dependency problem. This issue was considered during the initial development of FLEX and it was determined that forced vectorization gave correct results. However, this did not carry over to the CRAY-2. The problem was easily corrected by processing each row of elements in two passes, using a do loop increment of 2 to ensure that two adjacent elements having common nodes would not be processed during the same pass through the computational loop. This guaranteed that no vector dependencies would occur if vectorization was forced.

A second difference worth noting between the CRAY-2 and earlier CRAYs is the speed of scalar fetches from scattered memory locations. This issue involves the tradeoff between increased storage versus CPU time in the design of efficient processing strategies, in particular, to optimize storage for very large problems. For machines with 2-4 million words of storage, it was considered wasteful to store often duplicated material parameters for each element in the model, i.e., bulk and shear moduli. It even seemed extravagant to store a single material number of each element, since for a model containing one million elements, there are often less than 20 unique sets of material properties. The approach adopted on earlier CRAYs was to develop a group of unique site profiles, condensed to minimum storage, and a pointer table defining the site profile for each row of elements. In this way material properties of a homogeneous 1000×1000 element model can be completely defined by 1005 words of storage compared with 1,000,002 words if storing the material model number for each element and 2,000,000 words if storing the bulk and shear moduli for each element.

When computing the internal resisting forces for a row of elements, the site profile for the row is checked against the previous row's profile, and if they differ, the condensed site profile for the new row is expanded and loaded, with one word for each element. The element material properties are loaded into the bulk and shear moduli arrays in a scalar loop since this process does not vectorize on the CRAY-2. On earlier CRAYs, the overhead associated with retrieving the two material parameters at arbitrary locations in memory and placing them in consecutive locations in a local array (to facilitate vectorization) was a factor of two in the worst case, i.e., 2 times the homogeneous site performance. Although expensive, this approach significantly increased the size of models resident in memory, thus eliminating expensive disk I/O. In contrast, on the CRAY-2 it was found to increase the overall processing cost by at least a factor of 3.3.

The increased overhead follows because scalar fetches on the CRAY-2 are only slightly faster than on the CRAY-1, but since vector processing is about 2.6 times faster the scalar overhead is more significant. This can decrease the overall performance of the CRAY-2 to below CRAY-1 levels. After review, system analysts considered this scalar performance as normal for the CRAY-2's memory architecture and recommended replacing the single scalar loop fetching material properties with a series of gather calls. The result of this effort was to slow processing another 10 percent.

Based on these findings, it is clear that savings in storage afforded by using the site profile approach are irrelevant on the CRAY-2 with its large memory, in view of the significant CPU overhead incurred by scalar fetches of material properties. Consequently, the CRAY-2 version of FLEX has been modified to store both the bulk and shear moduli for each element in the model.

3.2. Performance comparisons

With the changes described above, FLEX achieves performance levels on the CRAY-2 that are about 2.6 times those on the CRAY-1 and 1.8 times the CRAY X-MP. These comparisons are for the same model and site profile description. For 2-D elastic axisymmetric problems, FLEX can compute 1.4 million element-timesteps per second of CRAY-2 CPU time, and for 3-D problems, it computes 0.28 million element-timesteps per second. For example, integrating a 1000×1000 element 2-D problem 2000 timesteps (2 billion element-timesteps) requires $2 \times 10^9 / 1.4 \times 10^6 \approx 1429$ seconds ≈ 24 minutes, and integrating a $100 \times 100 \times 100$ element 3-D problem 200 timesteps (200 million element-timesteps) requires $2 \times 10^8 / 0.28 \times 10^6 \approx 714$ seconds ≈ 12 minutes.

The major advantage of the CRAY-2 is of course its large memory. For the explicit finite element algorithm used in FLEX, memory requirements are 5 and 7 numbers per node in 2-D and 3-D respectively. Storage of synthetic seismograms typically adds another 10 to 30 percent to this, depending on the number of outputs points and components. Thus, the million element 2-D and 3-D examples cited above would require on the order of 6 and 9 million words respectively for the model and output. The calculations described below routinely use 10–30 million words, with no disk access. No degradation in performance is observed as memory usage increases, which has been verified for problems accessing up to 200 million words.

4. Large-scale 2-D simulations of a Maine refraction line

The first example calculation models a 200 km refraction line shot along an approximate southeasterly track across the state of Maine. This was the longest line instrumented by the United States Geological Survey during a 1984 cooperative refraction experiment in the north-eastern United States and southeastern Canada, e.g., see [4]. The line is perpendicular to the dominant structural features in the area and is reasonably approximated by a 2-D model. There were three large, high explosive shots recorded on the line—at each end and in the middle—but only the northwestern shot is modeled here.

4.1. The discrete model

The geologic model is illustrated in Fig. 1, showing a 50×183 km section of the Earth's crust and upper mantle with piecewise homogeneous structure—based on USGS data and interpretations [5]. Near-surface seismic velocity features are represented in some detail in the section, as well as the dominant topographic features (although not apparent in the figure). The finite element model was composed of a uniform 1000×3660 Cartesian mesh (i.e., non-skewed) of 50×50 meter elements, for a total of 3.66 million elements (7.33 million nodal equations). Interfaces and free surface topography were resolved step-wise by the mesh as in conventional finite difference discretizations. Clearly, the CRAY-2 is essential for a model of this size since the entire calculation (model and data) resides in-core to eliminate costly memory swaps to disk.

Although the geology is primarily 2-D on the Maine refraction line, the use of an effective point source clearly makes the problem 3-D. Since 3-D modeling on the scale of this experiment is impractical on any modern supercomputer (requiring on the order of 1000 times more power




Fig. 1. A geologic model of the Earth's crust and upper mantle on a southwest section across the state of Maine, from [5]. P-wave speeds are indicated in each homogeneous 'layer'. Finer near-surface velocity structure and topography are included but not apparent in the drawing.

and memory than presently available), a 2-D approximation is necessary. This was accomplished by posing the problem as axisymmetric, with the source on the axis of symmetry, i.e., the left (northwest) boundary of the model. Axial symmetry is preferable to a plane strain approximation because it exhibits the proper radial divergence, however, the resulting global model in 3-D, obtained by sweeping the geologic section about the axis, certainly appears unphysical.

Boundary conditions applied to the finite element model include symmetry on the left side as described above, a so-called absorbing or radiation condition on the bottom and right, and a free-surface condition on the top. The absorbing boundary condition used here and in subsequent finite element calculations described below is the simple normal impedance condition based on the unidirectional wave solution, $[\sigma] = -\rho c[v]$, relating jumps in stress and particle velocity. This is implemented for both P- and S-wave incidence and does a surprisingly good job for non-normally incident waves, particularly the P-wave. It is also readily vectorizable.

The seismic source used in the USGS refraction experiments was a 2000 pound cylindrical charge of chemical high explosive. The six inch borehole was 150 feet deep, filled to a depth of 75 feet with explosive, and topped off with rained sand. Considering that the entire borehole is contained in a single element of the large-scale finite element model, it was not practical to include details of the source directly. Instead a simple surface pressure, actually lumped vertical forces, with Gaussian distribution over a few nodes was applied to excite seismic waves. No attempt was made to reproduce the actual source's radiation pattern over takeoff angle. Linear and nonlinear source simulations will be considered further in a subsequent section.

The source time function is chosen to match the observed or expected frequency content at some range—in order to better compare synthetics to actual seismograms. However, for linear models it is often more useful to calculate the impulse or step response, from which synthetic seismograms due to an arbitrary source time function are obtained by convolution. The frequency content of the derived response is of course limited by discretization, which, for the low-order elements (linear interpolation) used here, typically means significant dispersion when a wavelength is supported by eight elements or less. This behavior is implicit in the step or impulse response.

4.2. Finite element synthetic seismograms

Calculated response of the finite element model is illustrated in Figs. 2 and 3, showing synthetic vertical velocity seismograms at two kilometer station increments between 0 and 180 km from the northwestern source. They show true ground motion, i.e., no instruments response is included. Step response is plotted in the bottom suite of seismograms, and after convolution with 2 Hz and 0.5 Hz wavelets in the middle and top suites respectively. The synthetics in Fig. 2 are scaled by range (multiplied by radius) to remove spherical spreading attenuation, and further multiplied by ten to accentuate early arriving body waves. The synthetics in Fig. 3 are scaled by the square root of range to remove the cylindrical spreading attenuation of surface waves and clearly show the Rayleigh wave.

The seismograms in Fig. 2 show a variety of body wave behavior, including reflection and refraction arrivals from the intermediate layers (P_g), reflections from the deeper interfaces (P^* and P_mP), mode conversion of Moho reflected body waves (P_mP) to Rayleigh waves by




Fig. 2. Suites of magnified vertical velocity seismograms at 2 km station spacing across the Maine model. The records clearly separate body and surface wave phases, which are labeled in the lower suite. The wavelet responses are obtained from the step response by convolution.




Fig. 3. Suites of unmagnified vertical velocity seismograms across the Maine model, showing predominance of the Rayleigh surface wave (R_g).

topography, and layer resonance of Moho reflections at distinct stations. They also show clear Rayleigh to Rayleigh reflections from lateral velocity variations in the top layer, as well as from topography. Observe that the various coherent phases are readily identified from the suites of seismograms by their slope, i.e., inverse phase velocity. These have been measured from the 2 Hz suite in Fig. 2 for the following interpretations.

The first arrivals out to about 110 km are P-wave reflections and refractions from the first interface, while from 70 km and beyond the arrivals transition into deeper reflections and refractions. Corresponding S-wave arrivals can also be seen preceding the Rayleigh wave, which is clipped in these synthetics. Discernible from 50 km (subcritical) and stronger after 100 km (supercritical) is the Moho reflection, which crosses the other body wave arrivals at about 170 km. At a range of 128 km a Rayleigh wave is seen to emerge at 24 seconds, traveling most strongly to the left but also to the right. The mechanism is identified as scattering of the Moho reflected P-wave by a 0.4 km step in topography (modeled by eight elements).

Figure 3 shows unclipped Rayleigh wave (R_g) arrivals across the model. The wave propagates about 2.65 km/s and travels halfway across the model when the calculation is terminated at 35 s.

The mesh is adequate to support 3–4 Hz fundamental mode Rayleigh waves so the observed dispersion of the 2 Hz wavelet (containing frequencies above 4 Hz) is partly due to numerical dispersion in addition to the more significant effects of lateral and vertical velocity variations near the surface. This figure clearly shows the relative strength of the Rayleigh wave compared to body waves for a surface pressure source, but hides all details. Of more use is the magnified response in Fig. 2. The middle synthetics show Rayleigh wave scattering by transitions in topography occurring between 12 to 25 km in range, and by lateral changes in wave speed at 33 and 55 km. Since the model's left side is assumed to be an axis of symmetry, apparent Rayleigh wave reflections are also seen. These are from the symmetric scattering features on the other side of the axis. This is of course a drawback of the axisymmetric approximation, however, the nonphysical reflected phases are easily identified. In general, this large-scale 2-D simulation produces a variety of propagation phenomena that would never be seen in a classical layered half-space analysis but is observed in real data.

5. Large-scale 3-D simulations of a scattering experiment

The second example calculation considered here models some rudimentary aspects of 3-D scattering experiments, performed by Reinke and Stump [6], in a small-scale alluvium site using five pound buried charges with seismometers and accelerometers placed around the shot out to 20–30 meters. One purpose of the experiments was to investigate azimuthal variability of ground motions by careful instrument calibration, exploration of site inhomogeneities, and the use of repeated axisymmetric shots. The principal question of interest here is the scattering produced by caliche lenses (cemented deposits) in the near-surface alluvial layer, and in particular, the interaction of characteristic lengths, i.e., wavelength, layer depth and size of the scattering feature. The following calculation was conducted to examine the feasibility of 3-D calculations to evaluate scattering by cemented lenses.

5.1. The discrete model

The geologic model for this case, illustrated in Fig. 4, is a $30 \times 30 \times 6$ meter quadrant of the explosive testbed, including two layers of alluvial material—a 3 meter low-speed surface layer over more competent alluvium. The finite element model is composed of $200 \times 200 \times 30$ finite elements, for a total of 1.2 million elements or 3.76 million nodal equations of motion. The model is truncated on the bottom by an absorbing boundary condition, and similarly on the outer boundaries, while the inner boundaries are symmetry planes. Elements in the lower layer are elongated in depth by a factor of two.

In real geologies, there are typically many scatterers with assorted shapes and orientations over the field of interest. For the purposes of this calculation, i.e., to determine propagation and source parameters for 3-D scattering simulations, a single ellipsoidal lens was modeled in the top layer, with nominal dimensions of $1 \times 3 \times 5$ m and oriented as in Fig. 4. The P-wave speed ratio between the caliche lens material and surrounding alluvium was taken as 2.6, based on field measurements. No attempt was made to smooth the lens boundaries using skewed elements, although comparisons of rough versus smooth scatterers should be made in future simulations.

The source in the field experiments is a five pound chemical explosive at the bottom of a two




Fig. 4. The 3-D finite element model of a $30 \times 30 \times 6$ meter quadrant of the scattering experiment, from [6]. Details of the source cavity and scatterer show actual discretization. The upper layer is shown transparent, and the lower layer discretization is drawn with only 75 elements on a side although 200 are used in the actual model.

meter borehole. Rather than simulate this source directly, a simple pressure function was applied to a source cavity obtained by voiding elements around the symmetry axis in the model, as illustrated in Fig. 4. The source's pressure history was a step function, from which responses to other source time functions are obtained by convolution, in the same manner as described above for the 2-D model.

5.2. Finite element synthetic seismograms

Three-component velocity time histories were recorded on two lines over the model's free surface. One extended from the source epicenter over the lens to the absorbing boundary (i.e., on the line of symmetry), and the other was perpendicular to this line, extending from above the lens centroid to the boundary. Vertical velocity synthetics are plotted in Fig. 5 on the radial and crossing output lines, with the radial line data scaled by range and the cross line data merely magnified.

These records show weak body wave phases followed by fairly strong surface waves. Some limited observations can be made from these results, e.g., on the radial line the lens is seen to scatter surface waves into body waves downstream, but interpretation is complicated by the




Fig. 5. Vertical velocity seismograms recorded on the cross and radial lines indicated by dots on the surface in Fig. 5. In-plane components of velocity were also recorded for polarization studies.

superposition of incident and scattered fields in the simulation. To overcome this it is necessary to run an additional free-field CRAY-2 simulation of the same geological model, i.e. with the scattering obstacle replaced by the surrounding medium. Subtracting seismograms from the two calculations yields the scattered field of interest. Of course this doubles the cost of analysis, however, for scattering studies it is essential for quantitative interpretation to decompose the incident and scattered wave fields.

Since the source geometry was rather coarse in this model, an additional free-field calculation would also be useful in evaluating azimuthal uniformity of the radiation pattern. In addition, a 2-D axisymmetric, free-field simulation could be used as a basis for comparison. Therefore, with the caveat that an additional free-field, 3-D simulation is required and a 2-D axisymmetric simulation is useful, these results indicate that this type of 3-D geophysical modeling is sufficient to produce useful data for experiment interpretation and planning. In order to draw further conclusions from the synthetic data in Fig. 5 it is necessary to perform further simulations. These are currently being done.

6. Seismic source modeling

One drawback of the models considered above is their inability to represent details of the source, since in large-scale simulations this region is only covered by a few elements. A solution is to either refine the mesh in the source region or perform a local source simulation and couple it to the global mesh. The latter approach is examined here, namely, the use of separate source simulations in highly refined models of the source region. Both linear and nonlinear models are considered.

Two types of source problems are modeled. The first includes 'nearly linear' surface or down-hole sources that apply vibratory or impulsive loads. These can be reasonably approximated by prescribing forces or velocities at an interface, with linear constitutive behavior assumed in the medium. The second type is the explosive source, for which strongly nonlinear material behavior dominates the near-source response. The explosive source involves a significant modeling effort to incorporate the proper energy release and coupling, in addition to difficulties associated with nonlinear constitutive behavior.

6.1. Elastic source modeling

A few examples of near-field seismic source models are examined here, assuming elastic properties. The idea is to model soil geometry around the source (mechanical vibrator, airgun, etc.) in some detail, and apply suitable pressures or velocities at the interface. The simplest geometry is the homogeneous half-space with prescribed surface pressure distribution and time function. A slightly more complex geometry is the borehole on the axis of an axisymmetric half-space, with either normal or shear tractions applied to all or part of the hole's boundary.

Only simple wavelet time functions are considered, representing the dynamic part of a total load. This total typically includes a significant static component, e.g., due to weight of the mechanical vibrator or sonde, which is ignored here since dynamic signals are of principal interest. It is permissible to have tension in the wavelet pressure boundary loading since this would be biased to net compression by the neglected static component.

The first example is an axisymmetric pressure loading over a small area of the free surface of a homogeneous 100×100 element model. The time function is a wavelet with center frequency chosen so that about 20 elements support it. A Gaussian spatial distribution of lumped vertical forces is assumed near the model axis. A sequence of vector snapshots shows the resulting velocity wave field in Fig. 6. The left side is the axis of symmetry, the right and bottom sides are absorbing boundaries, and the top is the free surface. The Gaussian pressure contributes significant forces only on the leftmost three or four nodes. Note that the simple absorbing boundary is quite effective for the P-wave, but less so for the S-wave at shallower angles, as well as for the Rayleigh wave.

To quantify this refined source solution or incorporate it in larger models, a useful approach is to record velocity time histories on a surface surrounding the source in its far field. For convenience this is typically a Cartesian surface (rectangle or box) at 10–20 characteristic source dimensions. In Fig. 7, results from the present calculation show velocities at a takeoff angle increment of 15° . These are rotated to display radial and tangential motion. This approach yields a simple spatial characterization of the source function. By calculating the impulse or step response instead of a particular wavelet, other time functions can be obtained by convolution.

The remaining elastic examples considered here involve pressure or shear loading on a section of a borehole. Figure 8 shows a snapshot sequence for a 75 ft hole with a pressure wavelet applied over a 6 ft segment centered at a depth of 37 ft. Similarly, Fig. 9 shows a sequence for the same geometry with a shear wavelet applied over the segment. The loads in both examples transition from zero to full traction over one element, rather than tapering. This is too strong a discontinuity for a discrete model to comfortably support, and consequently some spurious local oscillations (so-called hourglassing) are excited. Hourglassing is automatically removed at the element level by using orthogonality of the element's model shapes, e.g., see [7].

The above results are useful in comparing different source types or in characterizing non-ideal sources for input to a ray tracing model. These data are also necessary in order to interpolate input to coarser models.

6.2. Inelastic source modeling

There are various approaches to simulating underground chemical explosions. One is to apply a pressure history to the boundary of a cavity defined at some suitable 'elastic' radius; another is




Fig. 6. A sequence of velocity vector snapshots showing P-, SV-, and Rayleigh waves emanating from a Gaussian surface pressure (on the top) centered on the axis of symmetry (left side). The right and bottom boundaries have radiation conditions. The sequence starts at 150 time steps with a 50 step increment.




Fig. 7. Velocity time histories versus takeoff angle for the calculation in Fig. 6. Amplitude is normalized to a unit radius on radial lines from the center of the source region at 15° increments. The solid and dashed curves show radial and tangential motion, respectively, and quantify the source's radiation pattern.

to apply suitable initial conditions on velocity and pressure over an effective source region; and still another is to model the expanding gases from the source starting from the initial explosive volume. Of course the level of modelling detail can be further increased to include explosive dynamics, vaporization, etc. More rigorous modeling is seldom necessary unless local effects like cavity growth or cratering are important.




Fig. 8. A vector velocity snapshot sequence for a pressure wavelet applied to the sides of a 75 ft borehole at its mid-depth. The model is 200×200 ft. The pressure discontinuity causes significant spurious oscillation on the cavity boundary.




Fig. 9. A vector velocity snapshot sequence for a shear wavelet applied to the sides of a 75 ft borehole at its mid-depth. The model is the same as that for Fig. 8. The shear discontinuity is seen to cause much less spurious oscillation than the pressure discontinuity.

If no information is available on nonlinear constitutive behavior of the near-source medium, then the only alternative is to prescribe pressure or velocity on an 'elastic' cavity boundary. Details of the time function are either chosen to populate an observed frequency spectrum, or determined from experimental data in similar media. If a nonlinear model is available then bulk initial values can be used to incorporate additional source physics in the simulation. The nonlinear basis for the present source simulation is the cap model.

Chemical explosives typically exhibit confined detonation pressures of 2.9×10^6 psi (200 kbars) or less. Cap parameters are defined for pressures up to 200,000 psi in the alluvial-type media of interest here. Therefore, the cap model cannot represent material behavior in the immediate neighborhood of the detonated explosive. It is necessary to determine a radius where pressure has decayed to 200,000 psi and prescribe initial values over this volume determined from shock jump conditions. This is the basis for the bulk initial value approach used here. It has




Fig. 10. Suites of vertical velocity seismograms for the nonlinear and linear explosive source simulations. Upward velocity is plotted to the left. For ranges less than 85 ft the response is dominated by surface spall. Beyond this range nearly linear response is observed, yielding 30 Hz single cycle wavelet-type motion.

been used extensively by Sandler [8] and validated against a variety of explosive source experiments.

The model examined here consists of a vertical cylinder of explosive 10 ft long and 0.25 ft in diameter, with the upper end buried at a depth of 10 ft in homogeneous alluvium. Assuming a nonlinear pressure decay proportional to $1/r^2$ (for cylindrical geometry) yields the 200,000 psi radius at about 0.5 ft. Therefore, the volume over which initial conditions are prescribed is 1 ft in diameter and 11 ft long. This volume of material is initially pressurized to $P_0 = 200,000$ psi, with initial radial velocity set to zero on the axis and increasing linearly to V_0 at $r = 0.5$ ft, where V_0 is the particle velocity determined from the jump condition for a shock pressure of P_0 and a shock speed based on the secant modulus.

The finite element model examined here covers a 150 ft deep by 200 ft wide axisymmetric region of homogeneous material (alluvium) discretized into 151×200 elements. On the model axis, elements are 0.5 ft wide so that initial conditions are prescribed over a single column eleven elements long. Figure 10 shows vertical velocity synthetic seismograms for both nonlinear and linear simulations. The two lower suites are magnifications of the nonlinear calculation, while the upper shows elastic response for a low-pressure wavelet applied to the cavity boundary.

Since the explosive source is very energetic, soil immediately above and to the side of the model axis is spalled, i.e., launched at high velocity with complete tension cutoff. This clearly violates the assumptions of small strain theory and no claim is made that this late time behavior is realistic. However, the waves of interest have propagated beyond the source region before mesh distortions and constitutive uncertainties invalidate the solution. In any event, at a range of 85–90 ft, vertical velocity has decayed to a small fraction of that over the shot. This behavior is shown in the lower suite and magnified by ten in the middle suite to show additional details. Upward velocity is drawn to the left in the individual seismograms. Referring to the middle suite, beyond 110 ft the sequence of phases are, first a P-wave with upward motion (loading) arriving at the elastic wave speed, then a stronger P-wave about 10 ms later with downward motion (unloading), and followed by an S-wave with upward motion again. This up-down-up motion is suggestive of the common velocity wavelet approximation used as source time function in geophysics. Integrated yields a 30 Hz up-down (sinusoidal) cycle in displacement for the present nonlinear case.

The elastic simulation shown in the upper suite confirms that the slope of the phases in the nonlinear synthetics correspond to elastic wave speeds of the medium. Center frequency of the wavelet is 300 Hz, which is nonphysical but chosen to maximize separation between the phases while minimizing grid dispersion. Clearly, the frequency is not high enough to separate the shear and surface wave since they differ by about 8% in phase velocity for such geologic media.

7. Discussion and conclusions

This paper has described a number of discrete wave modeling problems in geophysics, including the simulation of large-scale 2-D refraction experiments, small-scale 3-D scattering experiments, and axisymmetric seismic sources, both linear and nonlinear. All of these problems require numerical solution because some level of nonseparability or nonlinearity precludes the use of classical analysis techniques.

Since the finite element algorithms employed operate at the lowest level of numerical sophistication practical for the problems at hand, i.e., linear, Cartesian elements and explicit

(leapfrog) integration, very little attention has been paid to theoretical background here since it is well documented in the finite element and finite difference literatures. One exception is the cap model for implementing nonlinear constitutive behavior, and so some theoretical background is provided on this topic. The principal aims of the paper are to demonstrate the simulation and performance capabilities of the CRAY-2, possibly the fastest and largest general purpose machine available today, and to familiarize the theoretical analysis community with some of the practical modeling issues and problems encountered in geophysics. These issues include (1) the capacity for large-scale inhomogeneous models and their efficient implementation, (2) very fast processing with very large, fast memory, and (3) capability for nonlinear behavior, material attenuation, and radiation boundary conditions.

7.1. Large-scale modeling capacity

Regarding CRAY-2 machine capacity for large-scale simulation, we can comfortably assert that 2-D models are well-resolved and appear cost-effective for most practical linear propagation problems, however, 3-D models remain resource limited and are not cost-effective. Here the term cost-effective means that simulations cost much less than the physical experiments. The reason for the disparity between 2-D and 3-D capability is that, for explicit calculations, computational and memory requirements grow like n^{d+1} where n is the number of nodes in a representative direction, and $d+1$ includes the problem's spatial dimension (d) and time dimension. For example, if it is necessary to double a model's size or mesh resolution (halve the spacing), then resource requirements increase by a factor of 8 in 2-D and a factor of 16 in 3-D. This additional factor in 3-D is all too often prohibitive, not necessarily in terms of memory but more often in terms of execution time.

In terms of the 3-D scattering model described above, the frequency resolution (element size) is adequate but the symmetry conditions and depth are too restrictive to reproduce the full experiment. As a consequence the model's size needs to be effectively doubled. For this more realistic model the execution time would increase from 1.5 hours to 24 hours, while memory requirements would increase from 20 million words to about 160 million. For this grid a numerical simulation would probably cost twice as much as the experiment.

In the case of the 2-D refraction model considered above, the original discretization had twice the element size of the final mesh. Unfortunately, this coarse model would not support wavelet frequencies much above 1 Hz, and since the experiments indicated at least 4 Hz signal resolution the element size was halved as a compromise. Consequently, the execution time went from an original estimate of 1 hour to over 8 hours, and memory quadrupled from 5 million to 20 million words. However, this simulation still cost a small fraction of one of the original refraction experiments.

7.2. Speed and memory

The principal requirement for the large-scale simulations described in this paper is minimizing disk I/O by retaining the entire model in fast memory. Thus, very large memories are necessary and the 256 million words available on the CRAY-2 are adequate. However, there is a hardware mismatch between scalar fetch and vector processing speed, which can seriously degrade performance. Fortunately it can be circumvented by using less sophisticated inhomogeneous

model representations at the expense of memory. This puts a serious limitation on problem size when the memory limit is 2 megawords (early CRAY-1), but is insignificant with 256 megawords available. Experiments with very large problems requiring upwards of 200 megawords show that there is no degradation in performance as memory usage approached the limits of the machine.

The major issue in large-scale, explicit wave simulation is processing speed. In order to do practical 3-D modeling, at least an order of magnitude increase in speed is necessary, for the reasons cited above. Furthermore, this capability should not cost much more than present CPU charges (\$1000 to \$3000 per hour say), otherwise we could use the CRAY-2 as it stands and pay \$25,000 to \$50,000 per calculation. The point is to make 3-D calculations timely and affordable by increasing speed while holding cost constant. The digital computer's forty year history and present state shows that this is not an unreasonable expectation.

There are two avenues available to increase processing speed: one is to increase the central processor's clock rate and the other is to increase the number of processors. A factor of three increase in speed (faster clock, more efficient vector processors) can be reasonably expected in the next generation of CRAY machines, and by using the four available processors a factor of 10 could probably be achieved. Unfortunately, the cost would be prohibitive since these multiprocessor supercomputers are in reality multicomputers, essentially four computers in one, with the cost directly proportional to the number of processors used. The alternative is to use a dedicated multiprocessor machine with slower processors but many more of them. This solution of course depends on the ability to logically partition the problem, assign pieces to the processors, avoid memory conflicts, etc. These are currently research questions but will be addressed soon since machines are available and there exist large-scale problems demanding solutions. However, it is not clear that CRAY-type machines will provide the necessary hardware and software.

7.3. Nonlinear behavior, material attenuation, and boundary conditions

We gather here some comments on other technical issues illustrated by the source models described above. These linear and nonlinear models are local refinements of the source regions in large-scale simulations, and all are in homogeneous media. They were readily calculated on a minicomputer rather than a CRAY.

The ability to model nonlinear material behavior is critical near energetic sources such as explosions. This is not to say that mechanical devices do not involve some level of nonlinearity, only that it does not obviously dominate the solution. Note that the nonlinear behavior associated with medium to low stress levels has not been characterized to the extent that high dynamic stress levels have. The important feature of nonlinear modeling with the cap and similar plasticity algorithms is their extensive use of stress state tests and logical branching. This branching effectively prevents vectorization of the algorithm. Consequently, nonlinear geophysical simulations are not vectorizable on CRAY-type supercomputers, hence are quite slow. An alternative is to use parallelism to speed up the cap algorithm, but this is not yet feasible on CRAYs.

None of the models considered here include anelastic attenuation, i.e., frictional damping, in the material behavior. Although this feature is straightforward in implicit frequency domain simulations, it is difficult to incorporate in the time domain. An implementation is available in FLEX, using a standard linear solid algorithm tuned to give constant damping over a given range

of frequencies. However, the requisite memory and arithmetic operations exact a significant cost. Such ad hoc attenuation models should be used with caution when signal attenuation is an important feature of a time-domain simulation.

Regarding boundary conditions, it is clear that discrete wave simulation models must be truncated in space by suitable radiation boundary conditions. For the present examples, the lowest-order condition assuming normal incidence of plane P- and S-waves was generally adequate. Higher-order conditions valid for wider incidence angles are available [9], but in many cases are unnecessary since boundary reflections from the model's side can usually be identified in the synthetic seismograms.

Since virtually all theoretical work on time-domain absorbing boundaries assumes a homogeneous medium, any implementation is degraded when properties vary along the boundary, e.g., on the side. This depth variation is the rule in geophysics and should be accommodated in future theoretical work. Some care must also be exercised on the model's lower boundary because the typical increase of propagation speed with depth in geologic models causes waves to turn away from the bottom, becoming more grazing there and less efficiently absorbed. Finally, note that radiation conditions in nonlinear models present a host of new difficulties.

7.4. Conclusions

The calculations described here give an indication of our present ability to simulate large-scale, time-domain wave propagation in nonseparable or nonlinear models. They show that one of the most powerful supercomputers available today, the CRAY-2, can indeed perform very large, 2-D, elastic simulations efficiently, despite an imbalance between scalar memory access and vector processing speeds. Of course this assumes that the principal processing loops are fully vectorized. Explicit, linear wave solvers are well suited to vectorization, in contrast to nonlinear constitutive algorithms for example.

The calculations also indicate that, for 3-D problems commonly encountered in geophysics, still more speed is required. This is not to say that we cannot do useful 3-D simulations with the CRAY-2. In fact the above scattering example shows that some limited but interesting 3-D simulations in geophysics can finally be addressed at reasonable cost. What must be emphasized is that cost grows so quickly in the third dimension, that only processor speed increases by one, or better yet, two orders of magnitude, will push affordable model size beyond the minimum required for adequate temporal and spatial resolution. Multiprocessor (not multicomputer) hardware with distributed algorithms (e.g., Connection type machines), or vector-parallel machines with suitable compilers (e.g., Alliant type), provide another avenue to achieving the necessary speed increase.

Simulations of the type described here are the only means of investigating propagation phenomena involving both range and depth dependent (nonseparable) or nonlinear model properties. Other options like ray tracing or boundary integrals involve idealizations that are unacceptable for this broader class of earth models. Note that discrete geophysical simulation has, here and elsewhere, been used principally for the forward problem. As the above performance issues are addressed more fully by software and hardware, it will eventually become possible to implement discrete models for the inverse problem, i.e., in an optimization loop with the synthetic data generated from discrete trial models. This possibility is perhaps the ultimate goal of geophysical simulation research.

Acknowledgment

We gratefully acknowledge support for this research by the Air Force Geophysics Laboratory under contract F19628-84-C-0102, and computer time on the CRAY-2 generously provided by the University of Minnesota's Supercomputer Institute.

References

- [1] D.K. Vaughan, FLEX users guide, UG8298, Weidlinger Associates, Palo Alto, CA, 1983.
- [2] K.-J. Bathe and E.L. Wilson, *Numerical Methods in Finite Element Analysis* (Prentice-Hall, Englewood Cliffs, NJ, 1976).
- [3] I.S. Sandler and D. Rubin, An algorithm and a modular subroutine for the cap model, *Internat. J. Numer. Anal. Meth. Geomech.* 3 (1979) 173-186.
- [4] S.L. Klemperer and J.H. Luetgert, A comparison of reflection and refraction processing and interpretation methods applied to conventional refraction data from coastal Maine, *Bull. Seism. Soc. Amer.* 77 (1987) 614-630.
- [5] J.H. Luetgert, Personal communication, United States Geological Survey, Menlo Park, CA, 1987.
- [6] R.E. Reinke and B.W. Stump, Stochastic geologic effects on near-field ground motions in alluvium, Submitted for publication.
- [7] D.P. Flanagan and T. Belytschko, A uniform strain hexahedron and quadrilateral with orthogonal hourglass control, *Internat. J. Numer. Meth. Engrg.* 17 (1981) 679-706.
- [8] I.S. Sandler, Personal communication, Weidlinger Associates, New York, 1987.
- [9] B. Engquist and A. Majda, Radiation boundary conditions for acoustic and elastic wave calculations, *Comm. Pure Appl. Math.* 32 (1979) 312-358.

ADA 201531

REPORT DOCUMENTATION PAGE				Form Approved OMB No 0704-0186	
1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release: Distribution Unlimited		
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE					
4. PERFORMING ORGANIZATION REPORT NUMBER(S)			5. MONITORING ORGANIZATION REPORT NUMBER(S) AFGL-TR-88-0292		
6a. NAME OF PERFORMING ORGANIZATION Weidlinger Associates		6b. OFFICE SYMBOL (if applicable)	7a. NAME OF MONITORING ORGANIZATION Air Force Geophysics Laboratory (LWH)		
6c. ADDRESS (City, State, and ZIP Code) 620 Hansen Way, Palo Alto, CA 94304			7b. ADDRESS (City, State, and ZIP Code) Hanscom AFB, MA 01731-5000		
8a. NAME OF FUNDING/SPONSORING ORGANIZATION DARPA		8b. OFFICE SYMBOL (if applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER F19628-84-C-0102		
8c. ADDRESS (City, State, and ZIP Code)			10. SOURCE OF FUNDING NUMBERS		
			PROGRAM ELEMENT NO. 61102F	PROJECT NO. 2309	TASK NO. G2
					WORK UNIT ACCESSION NO. AR
11. TITLE (Include Security Classification) (U) LARGE-SCALE, EXPLICIT WAVE SIMULATIONS ON THE CRAY-2					
12. PERSONAL AUTHOR(S) G. L. Wojcik, D. K. Vaughan, M. Barenberg and J. Mould M. B. Hulit*					
13a. TYPE OF REPORT Reprint		13b. TIME COVERED FROM 1Jan87 TO 31Dec87		14. DATE OF REPORT (Year, Month, Day) 1988 October 11	
15. PAGE COUNT 23					
16. SUPPLEMENTARY NOTATION Reprinted from Applied Numerical Mathematics 4 (1988) 47-70 North Holland (*AFGL/LWH, Hanscom AFB, MA 01731)					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP	Scattering, Explicit Wave Simulations, 3-D Modeling		
19. ABSTRACT (Continue on reverse if necessary and identify by block number)					
<p>Most time-domain, wave modeling problems in geophysics are intractable by classical analysis methods, due principally to nonseparability and to a lesser extent material nonlinearity. Therefore discrete numerical solutions are often necessary for the simulation of realistic models. Applications in 2-D and 3-D geophysical modeling are the subject of this paper, particularly as solved on a CRAY-2 supercomputer. Implementation and performance differences between earlier CRAYs and the CRAY-2 are described, including the discrepancy between scalar fetch and vector processing speeds. Explicit finite element solvers are applied to applications involving 2-D simulation of a seismic refraction experiment across the state of Maine, 3-D simulation of near-source scattering experiments, and both linear and nonlinear axisymmetric source simulation. Results show that the CRAY-2 allows cost-effective 2-D simulations of truly large-scale models, but only begins to be effective in 3-D for models of interest in geophysics. The large memory (256 megawords) is adequate but a speed increase of at least an order of magnitude is necessary for cost-effective 3-D. True multiprocessor capability (rather than 'multi-computer') provides an alternative to raw speed but involves another set of difficulties.</p>					
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified		
22a. NAME OF RESPONSIBLE INDIVIDUAL JANET C. JOHNSTON			22b. TELEPHONE (Include Area Code) 617-377-3767		22c. OFFICE SYMBOL AFGL/LWH

08 10 18 049